CIS 553: Networked Systems

Intradomain Routing

February 22, 2021

Agenda

- Discovery
  - DNS ← NEXT
  - DHCP
- Intradomain Routing
  - Distance Vector
  - Link State

Name Resolution (iterative)

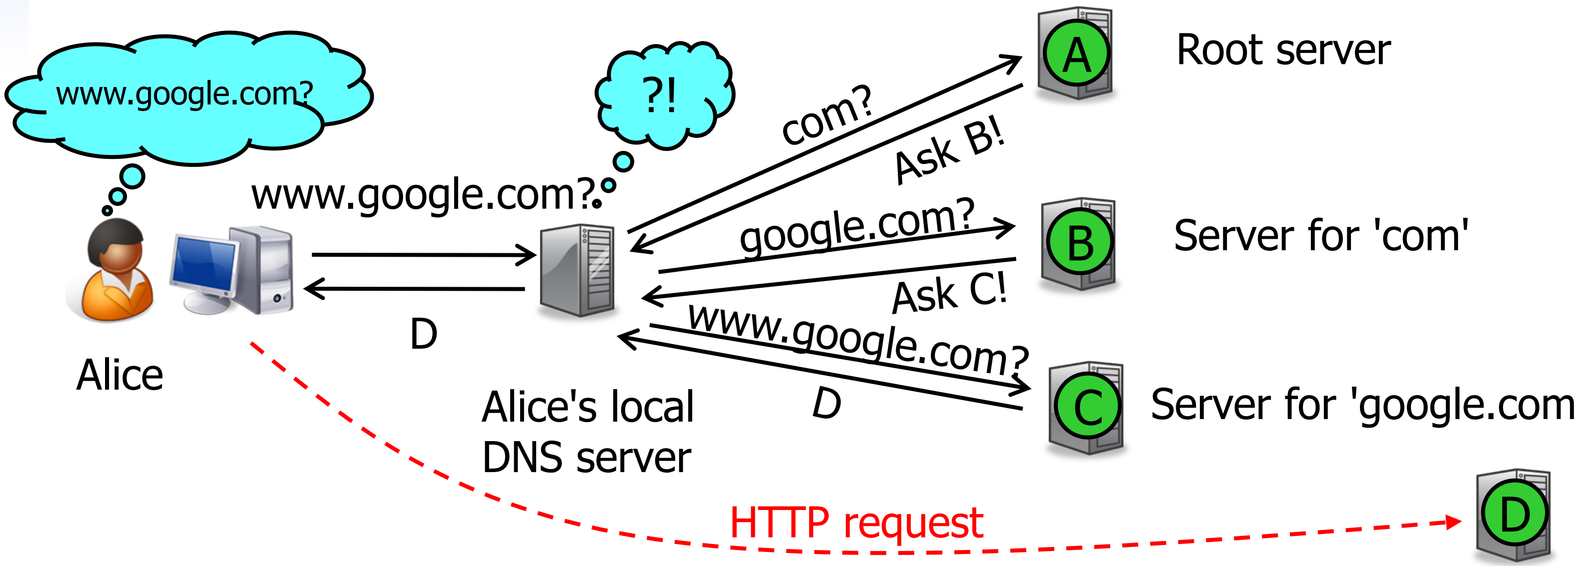

- Name lookup can be recursive or iterative
- Domain name is resolved step by step, starting with the TLD

Name Resolution (recursive)

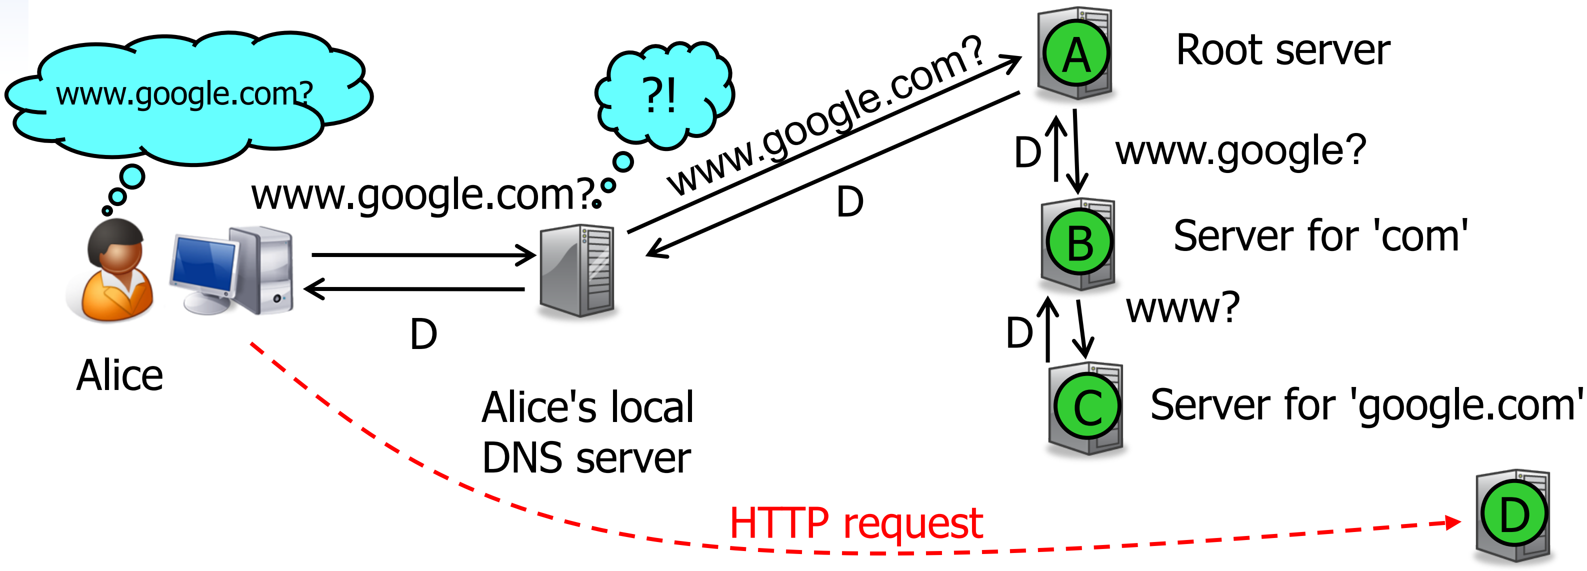

- Name lookup can be recursive or iterative
- Domain name is resolved step by step, starting with the TLD
- Recursive not used in practice. Why?

DNS Resource Records

DNS: distributed DB storing resource records (RR)

RR format: (name, value, type, ttl)

- Type: A
  - name is hostname
  - value is IP address
- Type: NS
  - name is domain (e.g. foo.com)
  - value is hostname of authoritative name server for this domain
- Type: CNAME
  - name is alias name for some "canonical" (the real) name

    www.ibm.com is really

    servereast.backup2.ibm.com
  - value is canonical name
- Type: MX
  - value is name of mailserver associated with name

Example lookup (simplified)

```
[liuv@carbon ~]$ dig seas.upenn.edu ANY
;; ANSWER SECTION:
seas.upenn.edu.   300   IN RRSIG MX 5 3 300 20180222185046 20180123175046 50475 upenn.edu.
JYRGUp8USVsmE7h4BI+kQ0kqB2qJOFf3T1/HQr3S/UXjCXK7D0pxTjG1 2OrF9TmFch5RRNUCj9z7xc9xYMZI+++Owo6KC5k4Sjv1+vYFd9K0YIp3
baeQTgZl5gTefOnV+3NCWd71SJBGv7de3r+fZnEMlwXe9bt/Uy7w83so s68=
seas.upenn.edu.   300   IN MX    10 telepathy.seas.upenn.edu.
seas.upenn.edu.   300   IN MX    10 psychopathy.seas.upenn.edu.
seas.upenn.edu.   300   IN MX    10 apathy.seas.upenn.edu.
seas.upenn.edu. 86400   IN RRSIG TXT 5 3 86400 20180207015318 20180108012245 50475 upenn.edu.
kQL+7PBRtw/K//KrIlsOt7JJdfr2DY2LAfktKGeWsSmeAnlrHKO5rnN+ RGyf64dcW3Z8jXEAGRdSKP+/+HJRheKE6GhJBdCd8fBEIapsi/3Aiyuu
8Z50O0f5bgoo/JVjl/BAJwdlV2s+sQjD3hviHnArvWWFgqDV+SDeJxhj ugk=
seas.upenn.edu. 86400   IN TXT   "v=spf1 ip4:158.130.64.0/21 ip6:2607:f470:8:64::/64 include:_spf.google.com
include:spf.protection.outlook.com include:spf-000c2a01.pphosted.com ~all"
seas.upenn.edu. 86400   IN TXT   "MS=576CC5A24F252658ACA66E552AD10E8015F6F153"
seas.upenn.edu.  3600   IN RRSIG NSEC 5 3 3600 20180222185046 20180123175046 50475 upenn.edu.
rWfhYBbLuVLyQFLHYukTVr9GJv4jeJnbDRR0vWbFUlT8T4XpuzUMSoWV H75DCL7FxrjiVu2DqsUpV7VUsd+xvyGh4lGlw8nudwSaCwgM+Y4XNnrt
dpnUrqH5O58kkv5bUs08ykPuRVpSyYl8+CXgZRGFlMAHh3s3bQ2nXWHM h94=
seas.upenn.edu.  3600   IN NSEC  080-STM-1.seas.upenn.edu. A MX TXT RRSIG NSEC
seas.upenn.edu. 86400   IN RRSIG A 5 3 86400 20180215172933 20180116165837 50475 upenn.edu.
IH+9CAn3cna2A9SWXsUYyI37R0BrrOzxUdPxuxZUnrN1tT1euWCwGQW4 NOERZbwYd+VzWmgJ5EKqA8NggGwePvZTLpzOqzlrD7d9jDP9X5UyXWH2
4FS+Flt/ZTpggo+gUACv2JfODRORiAhfJdV4Q95oFapyEqtCbVAsQCIX vwA=
seas.upenn.edu. 86400   IN A     158.130.68.91

;; AUTHORITY SECTION:
upenn.edu.      86400   IN NS    adns3.upenn.edu.
upenn.edu.      86400   IN NS    dns1.udel.edu.
upenn.edu.      86400   IN NS    adns1.upenn.edu.
...

;; ADDITIONAL SECTION:
dns1.udel.edu.  63029   IN A     128.175.13.16
adns1.upenn.edu.86400   IN A     128.91.3.128
adns3.upenn.edu.86400   IN A     128.91.251.33
...
```

Try it out: http://www.webdnstools.com/dnstools/dns-lookup

Goals: Are we there yet?

- Uniqueness: No naming conflicts
- Scalable
- Distributed, autonomous administration
- Highly available?
- Lookups are fast?

DNS Reliability

- Root servers are replicated
- Authoritative servers are also replicated
  - Parent can return multiple name servers
- When a request times out, retransmit
  - Exponential backoff when retrying same server
  - Can also try alternate servers
- Same identifier for all queries
  - Don’t care which server responds

Goals: Are we there yet?

- Uniqueness: No naming conflicts
- Scalable
- Distributed, autonomous administration
- Highly available
- Fast lookups?

DNS caching

- Performing all these queries takes time
  - Up to 1-second latency before starting download
  - Need to query every time
- Caching can greatly reduce overhead
  - The top-level servers very rarely change
  - Popular sites (e.g., www.cnn.com) visited often
  - Local DNS server often has the information cached
- Cached data periodically times out
  - Lifetime (TTL) of data controlled by owner of data
  - TTL passed with every record

Setting the Time To Live (TTL)

- TTL trade-offs
  - Small TTL: fast response to change
  - Large TTL: higher cache hit rate

- Following the hierarchy
  - Top of the hierarchy: days or weeks
  - Bottom of the hierarchy: seconds to hours

Negative caching

- Remember things that don't work
  - Misspellings like www.cnn.comm and www.cnnn.com
  - These can take a long time to fail the first time
  - Good to remember that they don't work so the failure takes less time the next time around
- Negative caching is now required
  - ~85-90% actually do it

Sneak Peak: DNS provides indirection

- Addresses can change underneath
  - Move www.cnn.com to 4.125.91.21
- Multiple names for the same address
  - E.g., many services (mail, www) on same machine
  - E.g., aliases like www.cnn.com and cnn.com
- Later: Name could map to multiple IP addresses!
  - Load-balancing
  - Reducing latency by picking nearby servers

Agenda

- Discovery
  - DNS
  - DHCP ← NEXT
- Intradomain Routing
  - Distance Vector
  - Link State

Bootstrapping Problem

- Host doesn't have an IP address yet
  - So, host doesn't know what source address to use
- Host doesn't know who to ask for an IP address
  - So, host doesn't know what destination address to use

- Solution: broadcast to discover a server who can help
  - Broadcast a DHCP server-discovery message

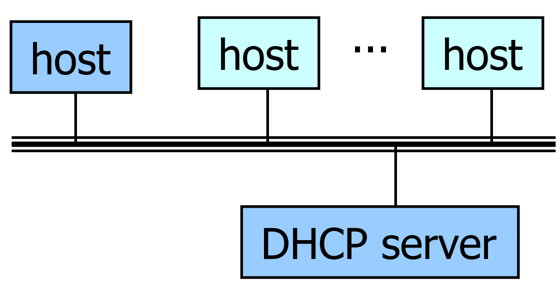

DHCP

- Dynamic Host Configuration Protocol
  - Defined in RFC 2131
- A host uses DHCP to discover
  - Its own IP address
  - IP address(es) for its local DNS server(s)
  - Basic routing information
    - IP address of gateway router
    - Prefix length of LAN

DHCP: Operation

- One or more local DHCP servers maintain required information
  - IP address pool, netmask, DNS servers, etc.
  - Application that listens on UDP port 67

DHCP: Operation

- One or more local DHCP servers maintain required information

Phase 1:

- Client broadcasts a DHCP discovery message
  - L2 broadcast, to MAC address FF:FF:FF:FF:FF:FF

DHCP: Operation

- One or more local DHCP servers maintain required information

Phase 1:

- Client broadcasts a DHCP discovery message
- One or more DHCP servers responds with a DHCP "offer" message
  - Proposed IP address for client, lease time
  - Other parameters

DHCP: Operation

- One or more local DHCP servers maintain required information

Phase 1:

- Client broadcasts a DHCP discovery message
- One or more DHCP servers responds with a DHCP "offer" message

Phase 2:

- Client broadcasts a DHCP request message
  - Specifies which offer it wants
  - Echoes accepted parameters
  - Other DHCP servers learn they were not chosen

DHCP: Operation

- One or more local DHCP servers maintain required information

Phase 1:

- Client broadcasts a DHCP discovery message
- One or more DHCP servers responds with a DHCP "offer" message

Phase 2:

- Client broadcasts a DHCP request message
- Selected DHCP server responds with an ACK

Agenda

- Discovery ✓
  - DNS ✓
  - DHCP ✓
- Intradomain Routing ← NEXT
  - Distance Vector
  - Link State

Recap: Forwarding

- Directing a packet to the correct interface so that it progresses to its destination
  - Read address from packet header
  - Search forwarding table

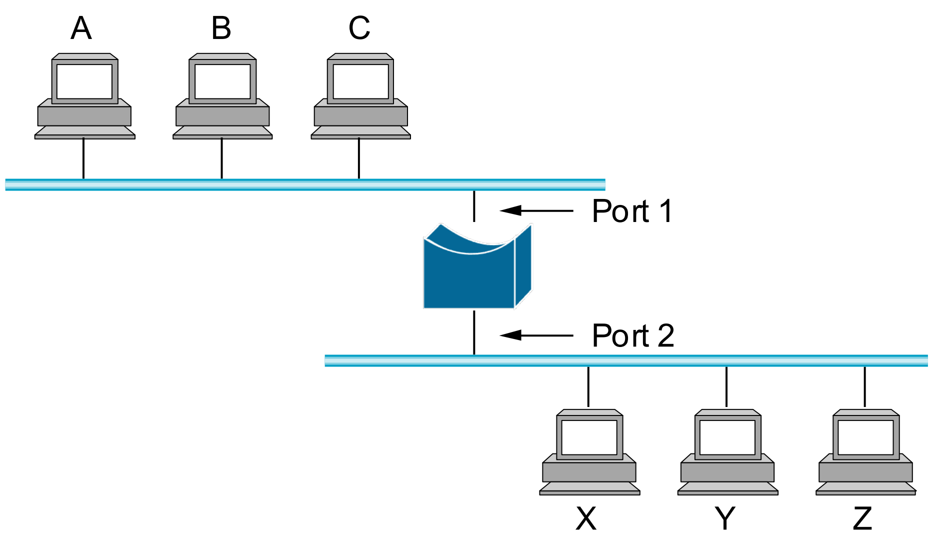

Host	Port
A	1
B	1
C	1
X	2
Y	2
Z	2

Forwarding table

Improving on the Spanning Tree

- Spanning tree provides basic connectivity
  - Wastes capacity
  - No guarantee of optimality
- Routing implies intelligence
  - Use all links to find "best" paths